

The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced?

Kevin Murphy^a, Rasmus M. Birn^a, Daniel A. Handwerker^a, Tyler B. Jones^a, Peter A. Bandettini^{a,b,*}

^a Section on Functional Imaging Methods, National Institute of Mental Health, Bethesda, MD, USA

^b Functional MRI Facility, National Institute of Mental Health, Bethesda, MD, USA

ARTICLE INFO

Article history:

Received 7 June 2008

Revised 10 September 2008

Accepted 17 September 2008

Available online 11 October 2008

ABSTRACT

Low-frequency fluctuations in fMRI signal have been used to map several consistent resting state networks in the brain. Using the posterior cingulate cortex as a seed region, functional connectivity analyses have found not only positive correlations in the default mode network but negative correlations in another resting state network related to attentional processes. The interpretation is that the human brain is intrinsically organized into dynamic, anti-correlated functional networks. Global variations of the BOLD signal are often considered nuisance effects and are commonly removed using a general linear model (GLM) technique. This global signal regression method has been shown to introduce negative activation measures in standard fMRI analyses. The topic of this paper is whether such a correction technique could be the cause of anti-correlated resting state networks in functional connectivity analyses. Here we show that, after global signal regression, correlation values to a seed voxel must sum to a negative value. Simulations also show that small phase differences between regions can lead to spurious negative correlation values. A combination breath holding and visual task demonstrates that the relative phase of global and local signals can affect connectivity measures and that, experimentally, global signal regression leads to bell-shaped correlation value distributions, centred on zero. Finally, analyses of negatively correlated networks in resting state data show that global signal regression is most likely the cause of anti-correlations. These results call into question the interpretation of negatively correlated regions in the brain when using global signal regression as an initial processing step.

© 2008 Published by Elsevier Inc.

Introduction

Spontaneous fluctuations in blood oxygenation level dependent (BOLD) fMRI signals have recently aroused a large amount of interest in the fMRI literature (Fox and Raichle, 2007). These fluctuations are often correlated between functionally related areas and can occur either on top of task-induced signal modulations (Fox et al., 2006) or in the absence of an explicit task (Biswal et al., 1995). It has been hypothesized that correlated fluctuations reflect synchronized variations in the neuronal activity of discrete brain areas and are representative of functional connections within networks of the brain. Functional connectivity analyses can investigate these coherent signal fluctuations, characterized by their low frequency (~0.1 Hz) (Lowe et al., 1998). By studying fluctuations at rest, researchers have claimed that the brain is intrinsically organized into dynamic, anti-correlated functional networks (Fox et al., 2005; Fransson, 2005; Greicius et al., 2003). The extent to which a commonly used correction method called global signal regression alters functional connectivity

maps by introducing anti-correlated time series is the topic of this paper.

Biswal et al. first reported the correlation in fMRI signal fluctuations between the left and right motor cortices when the brain was at “rest” (Biswal et al., 1995). Subsequent studies have identified several consistent resting state networks, including motor, auditory, visual, attention and default mode (Damoiseaux et al., 2006; De Luca et al., 2006; Greicius et al., 2003). The default mode network is of particular interest since it appears to be more active during rest than during task (Raichle et al., 2001). It has been hypothesized that this activation is indicative of internal monitoring, e.g. “day-dreaming” or a “wandering mind” (Buckner et al., 2008; Mason et al., 2007). Early clinical studies have attributed disruptions in the connections between nodes of the default mode network to disorders such as Alzheimer's disease (Greicius et al., 2004), schizophrenia (Garritty et al., 2007), ADHD (Sonuga-Barke and Castellanos, 2007) and autism (Just et al., 2007; Kennedy et al., 2006).

The common assumption in most fMRI investigations of connectivity is that correlated fluctuations in resting state networks are neuronal in origin. However, other sources of fluctuations exist in fMRI data that are not directly related to local neuronal firing and have a largely physiological source. Cardiac pulsations and respiration-related artefacts can cause significant correlated signal changes in

* Corresponding author. Laboratory of Brain and Cognition, National Institute of Mental Health, Building 10, Room 1D80B, 10 Center Dr. MSC 1148, Bethesda, MD 20892-1148, USA. Fax: +1 301 402 1370.

E-mail address: bandettini@nih.gov (P.A. Bandettini).

the vicinity of large blood vessels and throughout grey matter, obscuring spontaneous neuronal fluctuations (Lund et al., 2006). By measuring cardiac and respiratory traces during acquisition of fMRI data, the influence of the related physiological fluctuations can largely be removed, using retrospective correction techniques (Glover et al., 2000; Hu et al., 1995). Other techniques have been developed that do not require physiological recordings (Beall and Lowe, 2007; Chuang and Chen, 2001; Perlberg et al., 2007). However, further physiological sources of noise in the fMRI signal remain uncorrected by these methods. In particular, low frequency BOLD fluctuations have been associated with changes in the level of arterial carbon dioxide (CO_2) in the frequency range 0–0.05 Hz (Wise et al., 2004) resulting from changes in the breathed-volume (respiration volume over time – RVT) (Birn et al., 2006) and changes in the rate of cardiac pulsation at ~ 0.08 Hz (Shmueli et al., 2007). These sources of noise can artificially inflate connectivity measures since they introduce global, spatial coherence across the brain.

Global signal regression, otherwise known as orthogonalization to the global signal, is often performed as a processing step in an attempt to account for several potential sources of physiological noise. In this technique, the global signal, calculated by averaging the time series over all voxels in the brain, is used as a regressor in a general linear model (GLM) to remove the associated variance (Desjardins et al., 2001; Macey et al., 2004). This technique assumes that fMRI experiments are concerned with local changes in neuronal activity and that global signals represent uninteresting sources of noise. However, this assumption is only accurate when the global signal and experimental conditions are orthogonal to each other, that is, uncorrelated. The majority of the methodology literature regarding global signal regression examines its effect on task paradigms. Whether findings from such research are applicable to seed-based connectivity studies is unknown.

The global signal correction technique has been widely used in the PET imaging literature to remove global blood flow fluctuations with some debate as to the appropriate removal method (Andersson et al., 2001; Arndt et al., 1996; Fox et al., 1988; Friston et al., 1990; Ramsay et al., 1993). It has been suggested that defining the global signal as the average over all voxels can introduce a bias since the resulting time course will not be orthogonal to the task-induced activations (Andersson, 1997; Strother et al., 1995). Given that voxels responding to the experimental task are included in the global regressor, this correction technique will significantly influence the results by underestimating true activation levels and by introducing deactivations.

The global signal explains a large proportion of spatial variance in PET imaging (Fox and Mintun, 1989; Friston et al., 1990). Whether removal of this global signal is prudent in fMRI studies of task-related activity is a matter of debate. In spatially smoothed fMRI data, the global signal was shown to be strongly influenced by the performance of a behavioural task but inclusion of this signal as a covariate did not reduce sensitivity (Aguirre et al., 1997). Also, including the global signal covariate in a general linear model reduced spatial coherence and subsequently stabilised false-positive rates (Zarahn et al., 1997). However it has also been demonstrated that if the global signal is strongly correlated with experimental manipulations, considerably different results may be obtained. Aguirre et al. analysed an fMRI dataset in which subjects made button presses to brief intermittent stimuli both with and without covariation for the global signal and found that global signal regression reduced the spatial extent and intensity of positively activated regions whilst introducing spurious negatively activated areas (Aguirre et al., 1998a). Interestingly, these negative regions map out what is now known as the default mode network, a network shown to be anti-correlated with attentional tasks (i.e., deactivated during attention demanding tasks). Although the authors state that these areas “did not have a significant relationship with the task”, global signal regression may have helped bring these true deactivations to light.

Five different methods of global normalization in fMRI were compared by Gavrilescu et al., namely grand mean session scaling, proportional scaling, ANCOVA, a masking method and an orthogonalization method (Gavrilescu et al., 2002). The orthogonalization method, first proposed by Desjardins et al. (2001), was shown to perform better than other methods and is essentially the same as the global signal regression technique. However, it can, nonetheless, decrease the sensitivity of statistical analyses and induce artefactual task deactivations. Using both a high and low arousal emotional task, Junghofer et al. showed that the validity of proportional global signal scaling varied as a function of the emotional arousal of the stimuli (Junghofer et al., 2005). The high arousal condition violated the assumption of orthogonality between the global regressor and the experimental condition leading to a reduction in positively activated areas and “widely distributed artificial deactivations”, none of which made sense in terms of known emotional and default mode areas.

It is clear that global signal regression can cause reductions in sensitivity and can introduce false deactivations in studies of task activation since the assumption of orthogonality can be violated when the experimentally-induced activations contaminate the global signal. By definition, the experimental condition in resting state data is undefined. Although it is known that resting state fluctuations are low-frequency, the exact timing, spatial extent and relative phase between areas are unknown and may vary from session to session. The degree of correlation between the global signal and the resting state fluctuations cannot therefore be determined, and thus global signal regression could lead to spurious results in seed voxel correlation analyses of resting state data.

In a seed voxel analysis, a common method for computing functional connectivity, a time series is chosen in the brain that is hypothesized to represent fluctuations-of-interest. The correlation value between this time series and every other voxel indicates the extent of functional connectivity between the voxels. Using this connectivity analysis, previous studies have claimed that the brain is intrinsically organized into dynamic, anti-correlated networks (Fox et al., 2005; Fransson, 2005; Greicius et al., 2003; Kelly et al., 2008). These studies, however, have all used global signal regression as a pre-processing step. The default mode network is termed the task-negative network since it deactivates during a wide range of tasks. Its anti-correlated network is termed the task-positive network since it represents regions that increase in activity when default mode areas decrease. If the global signal is uncorrelated with the resting state fluctuations as represented by the seed voxel, then the finding of a task-positive network (anti-correlated with the task-negative network) at rest is valid. If, on the other hand, this is not the case, the interpretation that the brain is organized into anti-correlated networks may be brought into question since it is based on results that are contaminated by the global signal regression pre-processing step.

The extent to which global signal regression affects seed voxel functional connectivity analyses was investigated. This paper is organized into four sections (theory, simulations, breathing and visual data, and resting state data), each of which addresses a different aspect of resting state fluctuations and how they are influenced by global signal regression. The theory demonstrates that mathematically, regardless of the characteristics of the resting state fluctuations, global signal regression will always result in a negative mean correlation value during a seed voxel functional connectivity analysis. Three sets of simulations were performed. The first is an empirical demonstration of the theory. The second reveals that if similar resting state fluctuations (represented by a sine wave) are present in *all* voxels, functional connectivity analyses will determine that approximately half are anti-correlated with a seed after global signal regression. The relative phase of fluctuations in the seed voxel and other regions can bias those regions to become anti-correlated, serving as a possible explanation for the clustering of anti-correlated

areas in the brain. The third simulation investigated the behaviour of these anti-correlations as the spatial extent of voxels containing resting state fluctuations increases from zero voxels to the entire brain. The challenge in resting state functional connectivity studies is that the timing and location of spontaneous neuronal activation and global nuisance fluctuations (such as respiration) are unknown. Since it is impossible to systematically alter resting fluctuations, visual task activation was used in the breathing and visual data section to create localisable connectivity maps similar to those generated in resting state correlation studies. A comparatively global fluctuation was introduced into the data using a breath-holding challenge. The relative phase of the two tasks were varied to examine how the influence of global signal regression on connectivity measures is dependent upon the relationship between resting state fluctuations and the global signal, a quantity which, by definition, is unknown. Finally, task-negative and task-positive regions were defined in resting state data and changes in correlation measure with global signal regression were examined in depth.

Theory

This mathematical proof demonstrates that after global signal regression, the sum of correlation values with a seed voxel across the entire brain is less than or equal to zero.

Let each $S_i(t)$ be a column vector representing a i th voxel's time series, $i=1, \dots, N$. Since the mean of each voxel's time series can be removed during regression, assume that each voxel has a mean of zero. Let the corresponding time series after global signal regression be represented by the column vector $x_i(t)$.

Regression of the global signal is accomplished by solving

$$S_i(t) = g(t)\beta_i + x_i(t) \quad (1)$$

where the global signal (a column vector) is given by

$$g(t) = \frac{1}{N} \sum_{j=1}^N S_j(t). \quad (2)$$

Estimation of the regression coefficient, β_i , in Eq. (1) is given by

$$\beta_i = \left(g(t)^T g(t) \right)^{-1} g(t)^T S_i(t)$$

The average of the β_i coefficient across all voxels is given by

$$\frac{1}{N} \sum_{i=1}^N \beta_i = (g(t)^T g(t))^{-1} g(t)^T \frac{1}{N} \sum_{i=1}^N S_i(t) = (g(t)^T g(t))^{-1} g(t)^T g(t)$$

Thus,

$$\frac{1}{N} \sum_{i=1}^N \beta_i = 1 \quad (3)$$

Inserting Eq. (2) into Eq. (1) gives

$$S_i(t) = \beta_i \left(\frac{1}{N} \sum_{j=1}^N S_j(t) \right) + x_i(t). \quad (4)$$

Rearranging Eq. (4) provides the sum of the time series' after global signal regression:

$$\sum_{i=1}^N x_i(t) = \sum_{i=1}^N \left(S_i(t) - \beta_i \frac{1}{N} \sum_{j=1}^N S_j(t) \right) = \sum_{i=1}^N S_i(t) - \frac{1}{N} \sum_{i=1}^N \beta_i \left(\sum_{j=1}^N S_j(t) \right).$$

Inserting Eq. (3) gives

$$\sum_{i=1}^N x_i(t) = 0 \quad \forall t \quad (5)$$

Let x_1 be the seed voxel in a correlation analysis. Then the sum of dot products is given by

$$\sum_{j=2}^N x_1(t) \cdot x_j(t) = x_1(t) \cdot x_2(t) + x_1(t) \cdot x_3(t) + \dots + x_1(t) \cdot x_N(t). \quad (6)$$

From Eq. (5) we can determine that

$$x_2(t) = -x_1(t) - x_3(t) - x_4(t) - \dots - x_N(t).$$

Then Eq. (6) becomes

$$\begin{aligned} \sum_{j=2}^N x_1(t) \cdot x_j(t) &= x_1(t) \cdot [-x_1(t) - x_3(t) - x_4(t) - \dots - x_N(t)] + x_1(t) \cdot x_3(t) \\ &\quad + \dots + x_1(t) \cdot x_N(t) = -x_1(t) \cdot x_1(t). \end{aligned}$$

This implies that the sum of dot products is less than or equal to zero:

$$\sum_{j=2}^N x_1(t) \cdot x_j(t) \leq 0. \quad (7)$$

The sum of correlation coefficients is given by

$$\sum_{j=2}^N cc(x_1(t), x_j(t)) = \sum_{j=2}^N \frac{(x_1(t) - \mu_{x_1}) \cdot (x_j(t) - \mu_{x_j})}{N \sigma_{x_1} \sigma_{x_j}}$$

where μ_{x_i} is the mean and σ_{x_i} the standard deviation of the time series $x_i(t)$. Since we have assumed that $\mu_{x_i} = 0$ for $i = 1, \dots, N$, this equation reduces to

$$\sum_{j=2}^N cc(x_1(t), x_j(t)) = \sum_{j=2}^N \frac{x_1(t) \cdot x_j(t)}{M \sigma_{x_1} \sigma_{x_j}}$$

where M is the number of time points. If we assume that standard deviation values are unbiased across time series with negative and positive dot products, then the following equation holds:

$$\sum_{j=2}^N cc(x_1(t), x_j(t)) \leq 0. \quad (8)$$

Equation (7) demonstrates that if negative correlations don't exist before global signal regression, they must be introduced by the technique. For all voxels that correlate positively with the seed, negatively correlated voxels must exist to balance the equation. This is compatible with the previous suggestion that global signal regression causes correlation strengths to be distributed around zero (Buckner et al., 2008; Vincent et al., 2006), although the equations do not provide any information about the resulting distribution. This theory section demonstrates that in fMRI data, after global signal regression, a seed voxel correlation analysis must find anti-correlations, regardless of the seed voxel chosen or the signal fluctuations it represents. This fact is a clear problem when trying to interpret the meaning of negative correlations after using global regression.

Methods

Simulations

Simulations were performed to test the conclusions of the theory and to examine its practical implications. Two voxel and three voxel simulations were carried out using Matlab (MathWorks, Inc.). One thousand sets of time series were generated, each consisting of two and three time courses respectively. Resting state fluctuations were represented in each time series using a sine wave with a randomly chosen frequency between 0.1 Hz and 0.2 Hz and randomly chosen amplitude between 0 and 1. To this, Gaussian noise was added with zero mean and a standard deviation between 0 and 10. Each time

series lasted 1000 time points with an effective TR of 1 second. For each set of time series, the global signal was regressed using the GLM method. In order to verify Eqs. 3 and 5 empirically, the beta weights of the regression fits were averaged for each set and the sum of the time courses after global signal regression were calculated. The correlation coefficients between the first time course and the remaining one or two time courses in the set were calculated.

In another group of simulations, we investigated the effect of global signal regression on correlations between time series that exhibit similar resting state fluctuations that are slightly shifted in time. Null time series lasting 150 time points with a TR of 2 s were arranged on a 64×64 grid. A sine wave with a frequency of 0.1 Hz, amplitude of 1 and whose phase varied progressively along the x -axis from 0 to $\pi/4$ in equal increments was added to each time series to represent resting state fluctuations. Random Gaussian noise with a mean of 0 was added to each time series. The standard deviation of this noise was varied progressively along the y -axis from 0 to 10 in equal increments. The resulting dataset consisted of a 64×64 grid of time series with sine wave fluctuations of varying phase and increasingly large noise components. Global signal regression was performed by averaging over all 4096 time courses. Three correlation analyses were performed on the dataset, with and without global signal regression. The correlation of each time series with a sine wave of frequency 0.1 Hz and phase of zero was determined. Two 9×9 voxel regions-of-interest (ROIs) were drawn, one in a high and one in a low signal-to-noise ratio (SNR) region. The average time series in each ROI was determined and the correlation of the resulting time series with each voxel was calculated.

In the phase simulations above, every voxel contained a single frequency “resting state” fluctuation. A third “spatial extent” simulation investigated the influence of the fraction of voxels containing resting state fluctuations on correlation analyses when global signal regression is employed. One hundred artificial time series were generated lasting 300 s with a TR of 2 s. In each time series, physiological noise was simulated using a sine wave of 0.03 Hz (representing global signal changes, such as those induced by changes in depth of breathing) with random amplitude between 0 and 1 and random phase between 0 and $\pi/4$. Gaussian noise with a standard deviation of 1 was added. Resting state fluctuations were represented by a 0.1 Hz sine wave whose amplitude was 1 and 10 in low SNR and high SNR simulations respectively. This fluctuation was inserted in N of the 100 time courses, where N ranged from 1 to 99. For each iteration of the simulation, the global signal was calculated by averaging across the 100 time series and global signal regression was performed. A seed voxel time series was generated by regressing the same global signal from the 0.1 Hz sine wave resting state fluctuation. The average seed voxel correlation across the N voxels was determined along with the average across all remaining voxels. This whole process was repeated 100 times allowing study of the correlation values as a function of the percentage of voxels containing the fluctuations-of-interest.

Breath holding and visual data

This task combines both a breath hold block design and a visual block design task, varying the phase between the two in different conditions. The visual task produces a relatively small, focal activation in the visual cortex and was used as a model of resting state fluctuations. The global signal was modulated using a breath hold task that produces a large global BOLD signal increase (Kwong et al., 1992; Li et al., 1999). By combining the two, one can determine the effect of regression of a known global modulation on known focal “resting state” fluctuations (which are unknowns in actual resting state data). The extent to which the relative phase of the global and local signal changes influence local “resting state” connectivity measures was investigated.

Eight young healthy adults were scanned on a 3 T General Electric HDx MRI scanner equipped with a 16-element receive-only brain array coil under an IRB approved protocol after providing informed consent. Single shot, full k -space gradient recalled EPI was used for all functional scans. Whole brain coverage fMRI scans with 27 sagittal slices were collected with the following scan parameters: TR=2 s, TE=30 ms, matrix= 64×64 , FOV/slice=24 cm/4 mm, flip angle= 90° , reps=150. Physiological data were recorded during each scan using a pulse-oximeter placed on the left index finger and a pneumatic belt positioned at the level of the diaphragm.

The subjects experienced five conditions: *VisOnly*, *Synch*, *Synch+10s*, *Asynch* and *RandVis* (counterbalanced across subjects). In the *VisOnly* condition, an 8 Hz contrast reversing flashing checkerboard stimulus was presented using rear screen projection in a block design fashion 30 s OFF (grey fixation)/20 s ON. The subjects were asked to fixate on a cross in the centre of the screen. In the remaining conditions (*Synch*, *Synch+10s*, *Asynch* and *RandVis*), the subjects were asked to perform a breath holding task. This consisted of a 30 s countdown presented centrally with paced breathing instructions, that is, “Breath In” and “Breath Out”, each lasting 2 s. When the countdown reached zero, the subjects are asked to hold their breath for 20 s with a similarly presented countdown. The task was timed to present “Breath Out” directly before the “Hold Breath” instruction so that breath holding occurred after expiration. In the *Synch*, *Synch+10s* and *Asynch* conditions, a visual task identical to the *VisOnly* condition was presented together with the breath holding task with only the relative phase of the two tasks differing. In the *Synch* condition both the visual stimulus and breath hold onsets coincided temporally. In the *Synch+10s* condition, the visual task was delayed by 10 s with respect to the breath holding task. Since responses to breath holds are typically delayed compared to BOLD activation responses, the aim of this condition was to cause the peak BOLD response for each task to overlap temporally. In the *Asynch* condition, the visual task was presented so that its ON period ended when the ON period for the breath holding task commenced. Finally, in the *RandVis* condition, an 8 Hz visual checkerboard, presented in an event-related design with a variable inter-stimulus interval, was displayed along with the breath hold task (30 s normal breathing, 20 s breath hold) such that for each second long interval there was a 50% probability of checkerboard presentation.

Data were pre-processed using AFNI (Cox, 1996). Reconstructed images were first corrected for motion using a rigid-body volume registration. Physiological noise correction was performed using the RETROICOR technique to remove unwanted signal oscillations at aliased cardiac and respiration frequencies (Glover et al., 2000). Data were time shifted to align separate slices to the same temporal origin. Voxel time series were bandpass filtered between 0.01 Hz and 0.1 Hz and converted into percent change by removing and then dividing by the mean. Global signal regression, using all voxels within the brain, was carried out using a GLM. Datasets with and without the global signal removed were carried forward into the next analysis step.

In each subject, a functionally defined visual cortex mask was generated using the *VisOnly* condition (without global signal removed) by correlating each voxel's time series with the relevant task response and thresholding at $P < 1 \times 10^{-6}$. This task-induced response was modelled by convolving the task timing with a Gamma-variate haemodynamic response function ($\text{irf}(t) = k t^r e^{-t/b}$) where k is a scaling factor, $r=8.6$ and $b=0.547$; (Cohen, 1997). Similar correlation analyses were performed on all conditions, with and without global signal removed, using the relevant visual task model. Correlation values were averaged across visual cortex and then across subjects. A seed region correlation analysis was performed by determining the 10 highest t -statistic voxels in visual cortex in the *VisOnly* condition (without global signal removed). The seed region time series was generated by averaging over these voxels in each condition. The correlation coefficient between this time series and the time course in each voxel of the brain was determined. For each subject, a histogram of these correlation values for each condition was

calculated using voxels present in the visual cortex maps and also across all voxels in a whole brain mask. Group correlation maps were determined by converting each map to the standard stereotaxic coordinate system of Talairach and Tournoux (1988) using a 12-part linear transformation, averaging the arctanh of these maps and returning the tanh of the average (Fisher transformation).

Resting state data

Two resting state scans, each lasting 5 min, were collected on twelve young healthy adults with scanning parameters identical to the breath holding and visual data. Similar pre-processing steps were performed. After concatenation, the percentage change datasets, with and without global signal regression, were converted to the Talairach coordinate system as above. During this step, the data were resampled to 3 mm cubic voxels and spatially blurred using a Gaussian kernel with 6 mm full-width-at-half-maximum isotropic deviation.

Correlation maps were produced by extracting time courses from a seed region in the posterior cingulate/precuneus (PCC). This seed region was defined by drawing a 12 mm diameter sphere centered around the previously published Talairach coordinate $[-5, -49, 40]$ (Fox et al., 2005; Shulman et al., 1997). Time courses from voxels in this region were averaged and the correlation coefficient between this seed time course and each brain voxel was calculated. This process was carried out on the resting data with and without global signal removed. Group maps were obtained by averaging correlation values across subjects using a Fisher transform (for each voxel in Talairach space, the arctanh of the correlation value for each subject was averaged and the tanh of this average was returned). The individual subject correlation maps were thresholded at a correlation value of ± 0.35 corresponding to a P -value of 9×10^{-10} (which, when corrected for multiple comparisons using a Bonferroni correction, equals $P = 4.1 \times 10^{-5}$). From the maps on which global signal regression was performed, the task-negative network (i.e. the network that deactivates during a task and is more active during rest) was defined on a subject-by-subject basis as those voxels that passed threshold and displayed a positive correlation with the PCC ROI. Since the PCC lies within the task-negative network, every region that positively correlates with it also lies within that network. On the other hand, the task-positive network (i.e. the network anti-correlated with the task-negative network) was defined as those voxels passing threshold displaying negative correlations with the PCC.

The correlation values in each of the task-positive and task-negative networks were compared in both the resting datasets with and without global signal removed. Since global signal regression is hypothesized to eliminate uninteresting low-frequency fluctuations from the data, the correlation values in the task-positive and task-negative networks were compared with a method that purports to remove low-frequency BOLD fluctuations due to changes in breathing depth. Respiration volume per time (RVT) correction was performed by determining the envelope of the respiration trace, subtracting the minimum from the maximum, dividing by the period of the respiration and then removing this time course (and time-shifted replicas) from each voxel using multiple linear regression (see Birn et al. for more details, Birn et al., 2006). Correlation values were averaged across the task-negative regions for each subject for the three conditions: without global signal or RVT correction, with RVT correction and with global signal regression. Similar averaging was performed in the task positive regions. Histograms of the correlation values across the entire brain, with and without global signal regression, were calculated.

Average time series in both the task-positive and task-negative networks were calculated on a subject-by-subject basis for datasets with and without regression of the global signal. Correlations between all combinations of these time series and the global signal were determined. The phases between these time series before global signal

removal were calculated by determining the maximum correlation between both at multiple time lags. The percentages of voxels in the brain above threshold that positively and negatively correlated with the PCC region both before and after global signal removal were also calculated.

Results

Simulations

All 1000 sets of time courses in both the two voxel and the three voxel simulations demonstrate that Eq. (3) is true: the beta weights from the GLM average to exactly 1 ($\pm 10^{-15}$) in all cases. Each time point summed across all time series in the set after regression of the global signal is exactly zero ($\pm 10^{-13}$), demonstrating Eq. (5). The final correlation between the time courses after global signal removal in the two voxel case is exactly -1 ($\pm 10^{-15}$) in all 1000 simulations. In the three voxel case, the sum of correlation values is always less than zero with a mean of -1.02 and standard deviation of -0.05 over the 1000 simulations. The maximum value of the sum was -0.83 . This confirms Eq. (8) by demonstrating that after regression of the global signal, the sum of correlation values to a seed voxel is less than or equal to zero, regardless of the fluctuations it represents.

The phase simulation results are shown in Fig. 1. Correlations with the resting state sine wave demonstrate that the fluctuations are detectable in the majority of voxels even in low SNR time courses. Phase shifting the fluctuations by up to $\pi/4$ has little effect on detectability until global signal regression is performed, after which approximately half of the voxels display a negative correlation with the resting state sine wave. Correlations with averaged time series from ROIs show that the resting state fluctuations remain highly visible, whether the ROI is drawn from a high SNR (ROI1) or low SNR (ROI2) area, when global signal regression is not performed. After such a procedure, however, approximately half the voxels (52% and 48% for ROI1 and ROI2, respectively) display a negative correlation with the ROI. The phase simulations demonstrate that global signal regression can introduce negative correlations between time series with identical signal modulations that are slightly shifted in time.

The spatial extent simulations reveal that seed voxel correlation values are influenced by the size of the region containing the resting state fluctuations (see Fig. 2). Regressing the global signal from the data reduces correlation values as the spatial extent increases in both voxels containing fluctuations and voxels containing noise. Since noise voxels are not correlated with the seed voxel beforehand, global signal regression introduces negative correlations. In the low SNR case, this is a gradual decrease as spatial extent grows finally reaching a value of -0.3 . In the high SNR case, this reduction is more rapid with the correlation values in noise voxels reaching this value when only 10% of voxels contain fluctuations. Both SNR cases are extremes and the situation in real data is most likely to be somewhere between the two. These simulations demonstrate that the reduction of correlation values due to global signal regression is dependent on the number of voxels containing fluctuations with greater spatial extent leading to a greater decrease. Not only does global signal regression reduce correlations in areas containing fluctuations, it introduces negative correlations in pure noise voxels.

Breath holding and visual data

The breath holding and visual data demonstrate that global signal regression can severely affect the outcome of resting state correlation analyses, particularly when the global signal (as represented by breath holding) and the “resting state” fluctuations (as represented by visual activation in this study) are highly similar. When BOLD signal increases due to visual activation and breath holding overlap temporally, as in the *Synch+10s* condition (average correlation




Fig. 1. The results of the phase simulation are shown. Each panel displays correlation values on 64×64 voxel grid. The time series for each voxel consisted of a sine wave representing resting state fluctuations with a frequency of 0.1 Hz and a phase that varied for voxels progressively from 0 to $\pi/4$ along the x-axis. Random Gaussian noise with a mean of 0 and a standard deviation varying from 0 to 10 along the y-axis was added. The first column shows correlations to a 0.1 Hz sine wave with zero phase. The remaining columns depict correlation values to a time series averaged over the ROIs shown. This figure demonstrates that small phase differences between time series can lead to negative correlations after global signal regression.

between visual task regressor and global signal was 0.61 ± 0.10 , regression of the global signal severely reduces correlation values (see Fig. 3). In contrast, correlations are unaffected in the *VisOnly* condition since the global signal (which was not modulated by breath holding) is unrelated to the fluctuations-of-interest (correlation between task regressor and global signal: 0.08 ± 0.16). Similarly, when the global modulation (breath hold) is unrelated to the “resting state” fluctuations (visual activation), as in the *RandVis* condition, global signal regression only has a small impact on correlation values. The extent to

which the global signal and the “resting state” fluctuations in the seed region are similar (i.e., the correlation value between the visual and global task) predicts almost exactly ($r = 0.98$) how much the global signal regression method will change the average correlation values in voxels containing the fluctuations.

Histograms of seed region correlation values across the visual cortex and across the whole brain are shown in Fig. 4. In all conditions, the distribution of correlation values in the visual cortex is skewed towards high values as one would expect. Global signal regression has little effect on the distribution of visual cortex correlation values in the




Fig. 2. The spatial extent simulations show how correlation values vary after global signal regression as a function of the percentage of voxels that contain resting state fluctuations. Red lines represent the average correlation over voxels that contain a sine wave (representing a resting state oscillation) whereas blue lines depict the average correlation over the remaining noise voxels. Both low SNR (amplitude of sine wave equals standard deviation of noise) and high SNR (amplitude of sine wave 10 times greater) simulations are shown, in solid and dashed lines respectively. Significance lines showing $P < 0.001$ are depicted. These simulations demonstrate that the reduction of correlation values due to global signal regression is dependent on the number of voxels containing resting state fluctuations with greater spatial extent leading to a greater decrease. Voxels containing purely random noise become more negatively correlated as spatial extent increases.




Fig. 3. Correlation values averaged over the visual cortex and over subjects are shown for the combined breath holding and visual task. Each of the five conditions demonstrate how the relative phase of global (breath hold) and “resting state” (visual) signal changes affects the correlation measure with and without global signal regression (red and blue lines respectively). The significance levels for all comparisons were determined using paired *t*-tests: *VisOnly* $P = 0.42$, *Synch* $P = 0.23$, *Synch+10s* $P = 8.9 \times 10^{-6}$, *Asynch* $P = 7.7 \times 10^{-4}$, *RandVis* $P = 0.045$. When the phase of the breath hold and visual responses are the same (*Synch+10s* condition), global signal regression causes a large reduction in correlation value in the visual cortex.

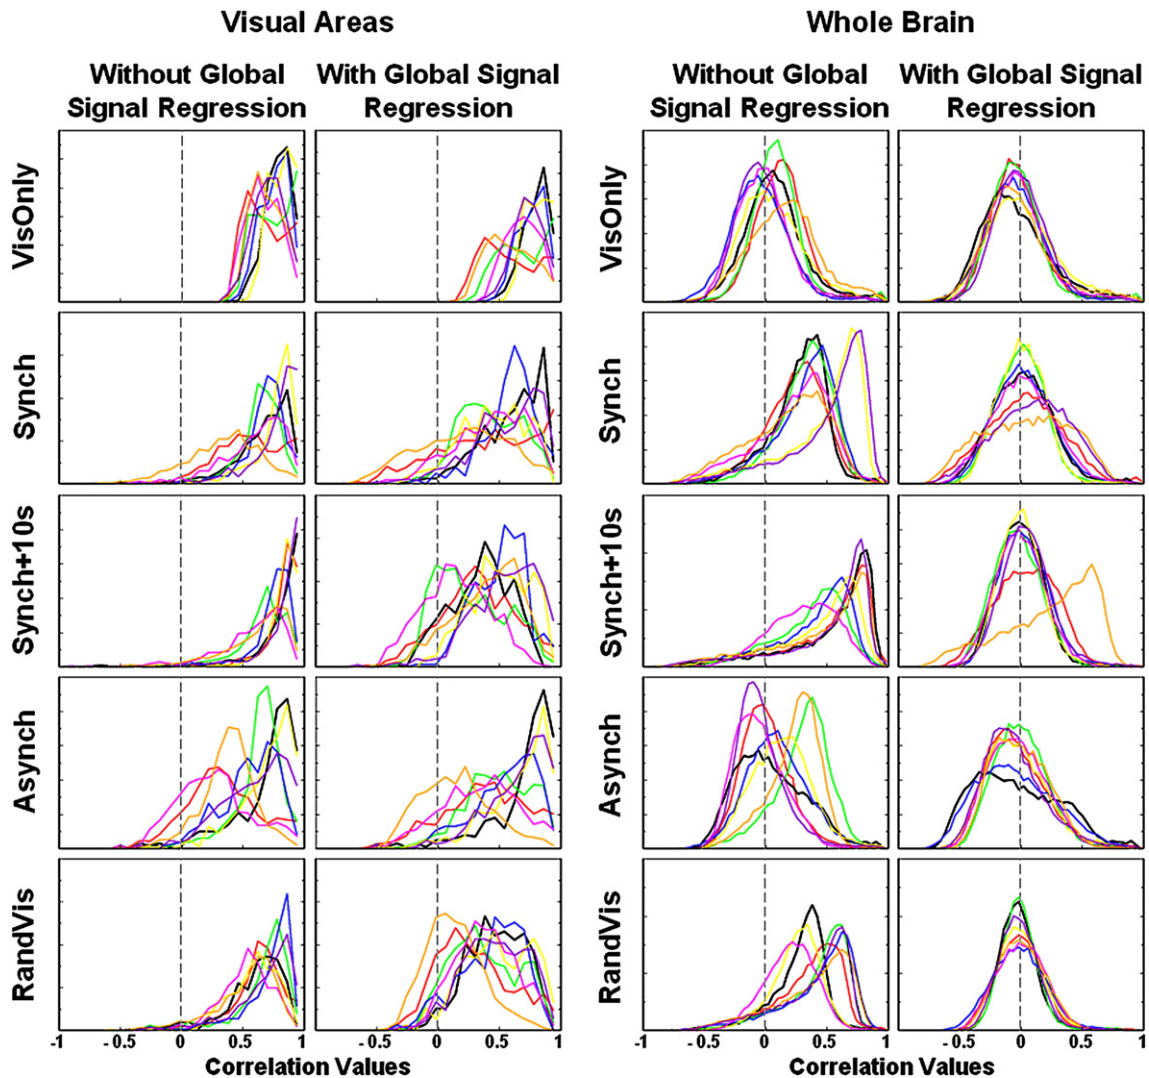


Fig. 4. The distributions of correlation values in the seed region correlation analyses are shown with each subject represented by a different colour. Each of the five rows displays one of the five breath holding and visual task conditions. Distributions drawn from only the visual cortex are displayed on the left with whole brain distributions on the right, for data with and without global signal regression. The distributions from the visual areas show that global signal regression is not revealing underlying neurally-related BOLD signal changes as we would like it to. Distributions in the visual cortex are broadened instead of sharpened by this technique and negative correlations are introduced. The whole brain distributions on the right demonstrate that global signal regression forces a bell-shaped curve, centred on zero, regardless of the relative phase of global and “resting state” changes.

VisOnly condition since the activation response and the global signal are unrelated. In the other conditions, global signal regression tends to widen the distribution towards lower values. Qualitatively, the greatest broadening of the distributions occurs in the *Synch+10s* condition when the global signal is most correlated with the visual task regressor representing resting state fluctuations (0.61 ± 0.10), forcing correlation values to extend into the negative range. If global signal regression properly removes global manipulations, all distributions should resemble the *VisOnly* condition after performing the technique.

The whole brain distributions show that the breathing task causes a global modulation that skews the distributions when no global signal regression is performed. In all conditions, independent of the relationship between the global signal and the “resting state” fluctuations, global signal regression causes these skewed distributions to become more bell-shaped and centered on zero, thus introducing negative correlations. Assuming that global signal regression removes purely respiration related changes to reveal true neuronal fluctuations, the seed region connectivity maps after global correction should be identical, regardless of the breathing manipulation. Fig. 5 shows that large differences exist in the resulting

correlation maps, none of which spatially resemble the *VisOnly* condition. Although voxels from the visual cortex are most positively correlated with the seed region in all conditions, areas that negatively correlate vary across breathing manipulations. The spatial organization of the negative voxels introduced by the correction is dependent upon the relative phase of the fluctuations-of-interest and the global signal. Coupled with the fact that highly correlated voxels remain highly correlated, this suggests that, after global signal regression, positive correlations may be more reliable than negative correlations. The extent to which these maps differ shows that global signal correction does not reliably remove the global artefact induced by the breath holding task.

Resting state data

Intrinsic correlations between a seed region in the PCC and all other voxels in the brain during a resting task are shown in Fig. 6. After global signal regression, approximately the same number of voxels are negatively correlated with the PCC seed region (task-positive network) as are positively correlated (task-negative network). When global signal regression is not performed, however, very few




Fig. 5. Three axial slices of correlations with the PCC region averaged across subjects are shown for each of the Breath Holding and Visual Task conditions. Without global signal regression, correlation values can be high across the entire brain. The *Asynch* condition shows that this is not the case when global signal and “resting-state” fluctuations are out of phase. If global signal regression performed the intended correction, one would expect all resulting maps to be similar. Large variation exists in the pattern of activated and deactivated regions across the conditions.

negatively correlated voxels pass threshold. The spatial distribution of the positive correlations is similar in both cases.

To probe the data more quantitatively, task-negative (correlated with the PCC) and task-positive (anti-correlated with the PCC) regions were defined for each subject. Average correlation values in these regions are shown for each subject in Fig. 7. Three conditions are depicted: the first draws values from data that have not been corrected for global fluctuations, the second from data that has undergone RVT correction and the third from data which has had the global signal regressed from it. RVT correction has been shown to remove some of the low-frequency fluctuations due to variations in depth of breath which causes spatially widespread confounds across grey matter and regions with large vessels including the PCC area (Birn et al., 2006). Unlike with global signal regression, correlations with the PCC did not greatly reduce when this correction was performed. Correlation values from the RVT corrected data are almost identical to those from the uncorrected data in both the task-negative and task-positive regions. On the other hand, global signal regression reduces correlation values in both areas. Reductions in the task-negative areas vary from 0.02 to 0.25 across all subjects. In task-positive areas, there is a greater variability across subjects in correlation values from the uncorrected data, ranging from -0.24 to

0.49 with a standard deviation of 0.23 . The majority of subjects display strong correlations with the PCC in task-positive areas when global signal regression is not performed. After removal of the global signal, task-positive areas are strongly anti-correlated with the PCC and variability in the correlation values across subjects is greatly reduced (average correlation value is -0.42 ± 0.07).

Histograms of the correlation values for each subject are shown in Fig. 8. The distribution is highly skewed towards positive values when global signal regression is not performed. Those subjects that show a negative correlation in task-positive areas in Fig. 7 correspond to the subjects with distribution peaks closest to zero. After global signal regression, the distributions revert to a more bell-shaped curve centred on zero. The tails of the distributions stretch further into the positive than into the negative correlation values. If correlation values in each voxel are compared on a voxel-by-voxel basis, before and after removal of the global signal, voxels on the high end of the distribution remain on the high end and voxels on the low end remain on the low end (data not shown). Qualitatively, comparing the distributions with those of Fig. 4, the resting state data appear to correspond most to the *Synch+10s* condition. In this condition, the global signal was most highly correlated with the fluctuations-of-interest. If there are indeed two populations of voxels that are highly anti-correlated with each

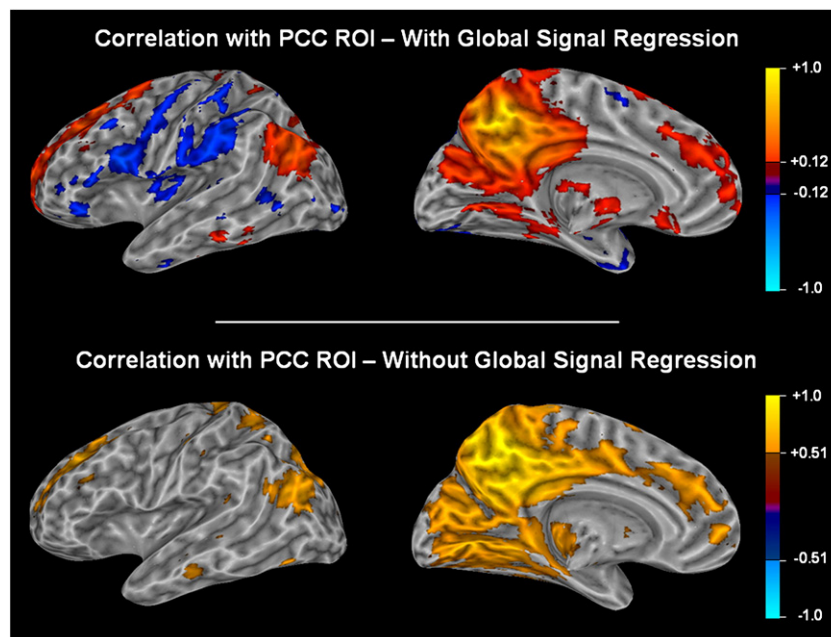


Fig. 6. Correlations to a seed region in the PCC averaged across all subjects using a Fisher transform are shown with and without correction for the global signal. The task-positive network is visible only after global signal regression. Since global signal regression changes the distribution of correlation values (see Fig. 8), it would be incorrect to use the same threshold for both maps. To enable a fair comparison, we chose a correlation value such that all voxels that positively correlate to the PCC and that passed threshold would comprise 15% of the total number of voxels. Using these thresholds, the default mode network is visible both with and without global signal regression, but negatively correlated voxels are only visible after the global correction. Fig. 7 and the distributions in Fig. 8 demonstrate that lowering the threshold in the map that retains the global signal would not reveal negatively correlated voxels in the task-positive network.

other, each including a large proportion of voxels in the brain, one would expect to see a bimodal-like distribution after global regression. A widened distribution in Subject 5 or the bump around -0.5 of Subject 9's distribution may be evidence for this. However, the majority of subjects demonstrate no such behaviour suggesting the bell-shaped distributions of the correlation values are an artefact of the mathematics of the global signal regression procedure which forces negative correlations to appear.

Table 1 shows the sizes of regions that correlate both positively and negatively with the PCC before and after global signal regression. When the global signal is included, areas with a significantly positive correlation can constitute as little as 25% of voxels up to almost the entire brain (96%). Negatively correlated voxels are almost non-existent with 0.59% being the greatest percentage in the twelve subjects. After global signal regression, positive correlations are reduced and negative correlations are increased (task-negative region = $9.65\% \pm 5.12\%$, task-positive region = $5.11\% \pm 4.65\%$). Phase differences between the task-negative and the task-positive time series before global correction range from 0.2 s up to 9 s. Comparisons with the global signal show that both the task-negative and task-positive time series correlate highly without removal of the global signal (0.75 ± 0.19 and 0.67 ± 0.18 respectively) and also display a large range of correlation values, from -0.51 to 0.62 . After global signal regression, task-negative and task-positive time courses become highly negatively correlated with an average of -0.71 ± 0.14 . The difference in correlation values before and after global signal regression spans a large range from 0.46 to 1.38.

The resting state data demonstrate that, when using the PCC as a seed region in a correlation analysis, anti-correlated networks are not significantly present until global signal regression is performed. As demonstrated in the breath holding and visual task, global signal regression renders the distribution of correlation values to be bell-shaped, centred on zero. After global signal regression, subject-to-subject variability in correlation values drawn from task-negative and task-positive areas is greatly reduced. Such regularized trends across subjects are a consequence of the global signal regression method.

Discussion

The simple fact that any two time series become perfectly anti-correlated when their global signal is regressed from each should give cause for concern when interpreting anti-correlations. If the first time series contains a signal modulation of interest, the second will be negatively correlated with the first after global signal regression, regardless of the signal fluctuations it contains. Extending this concept to multiple time series, the theory demonstrates that global signal regression changes voxel time series such that approximately half become negatively correlated with a seed voxel. Due to the circularity of the technique (i.e., deriving a nuisance regressor from the data itself), all voxels could have similar interesting signal fluctuations but will display a range of seed voxel correlation values from highly positive to highly negative after its implementation. These findings alone should cause fMRI researchers to question the efficacy of this technique as a correction tool.

Resting state data are unique in that the researcher does not have *a priori* knowledge of task signal modulations. The simulations and breath holding data yield an insight into the effect global signal regression has on known fluctuations. For example, the spatial extent simulations show that time series in noise voxels outside an ROI containing resting state fluctuations become negatively correlated with those fluctuations when global signal regression is implemented. Not only does global signal regression reduce correlations in areas containing fluctuations, it introduces negative correlations in pure noise voxels. The degree of negative correlation depends on the extent of the ROI and the signal-to-noise ratio of the fluctuations. These simulations demonstrate that if there are large scale resting state networks in the brain, their influence on the global signal will force purely noise voxels to become negatively correlated with the resting network (e.g. the PCC region). When the SNR of the fluctuations is suitably high, only $\sim 10\%$ of the brains voxels need contain these signal modulations for the remaining voxels to become negatively correlated. The information in Table 1 suggests that in resting state data the task-negative network is large enough for this situation to occur.




Fig. 7. Average correlation values across the task-negative (top panel) and task-positive (bottom panel) regions are shown for all subjects. Three conditions are depicted: 1) without any global signal or RVT correction, 2) with RVT correction to remove low-frequency fluctuations related to breathing depth and 3) with global signal regression. RVT correction has little impact on correlation values in both the task-negative and task-positive areas. Before global signal regression, task-positive areas can be highly correlated with the PCC region with only two subjects showing small negative correlations. Global signal regression causes all task-positive areas to become negatively correlated and reduces the variability across subjects.

The purpose of global signal regression is to remove any confounds that may mask neuronally-induced fluctuations. When confounds and resting state fluctuations are of similar frequency, as in the Breath Holding and Visual Data, their relative phases can severely affect the detectability of activation. In seed connectivity analyses where an explicit breathing confound exists, if global signal regression was completing the intended correction properly, one would expect the resulting connectivity measure distributions to resemble those of the *VisOnly* condition. However, this is not the case: global signal regression introduces negatively correlated voxels into the visual cortex. Even when no relationship between the global confounds and the fluctuations-of-interest exist (as in the *RandVis* condition), global signal regression does not sharpen the connectivity measure distributions as it should. These data demonstrate that global signal regression does not perform a correction that reveals the underlying neuronally-induced fluctuations.

In the Breath Holding and Visual Data, since the visual cortex represents such a small proportion of voxels it is expected that the seed voxels drawn from this area should bear little relation to most other voxels in the brain. When the breathing task introduces a global

confound, whole brain connectivity measure distributions are skewed. Removing the global confound with regression results in distributions that are bell-shaped. This is unsurprising since the time series resulting from global signal regression are the residuals from a GLM analysis that assumes a Gaussian noise model. However, the spatial pattern of activations and deactivations varies depending on the breathing manipulation (see Fig. 5). This indicates that global signal regression does not remove the global breathing confound successfully. Although the visual cortex remains the most positively correlated area (as it should), negatively correlated areas can be situated in vastly different regions after global signal regression, suggesting that, in this case, the negative correlations are not neuronal in origin but an artefact of the global signal regression method.

The task-positive network is not evident in resting state data until the global signal is regressed from each time series. This is not a thresholding artefact; Figs. 7 and 8 demonstrate that lowering the threshold in the uncorrected map would not reveal negatively correlated voxels in the task-positive network. If it is possible to see positive correlations to the PCC before global signal regression, why are the negative correlations not visible? The standard answer is that the uninteresting, confounding fluctuations which are removed by the technique overshadow the underlying neuronally-induced fluctuations. However, removal of low-frequency confounds due to breathing depth changes using RVT correction barely alters correlation values. Even if RVT correction does not remove all global signals of interest, it is a global confound that is known to overlap well with the default mode network and correlates significantly with the global signal at




Fig. 8. The distribution of correlation values to the PCC across all voxels in the brain for each subject are shown before (top panel) and after (bottom) global signal regression is performed. As with the results shown in Fig. 4, global signal regression causes the distributions to become bell-shaped and centred on zero.

Table 1
Descriptive statistics for resting state time series

Sub	Size before glob. sig. regr.		Size after glob. sig. regr.		Neg/Pos phase diff (s)	Corrs before glob. sig. regr.			Corrs after GSR	Difference in corr. values
	% Voxels with +ve PCC corr	% Voxels with -ve PCC corr	Task neg size (%)	Task pos size (%)		GS/Neg	GS/Pos	Neg/Pos		
S1	66.56	0.02	7.01	4.53	2.6	0.77	0.65	0.22	-0.58	0.80
S2	82.67	0.09	4.93	3.80	0.6	0.86	0.84	0.51	-0.79	1.29
S3	25.79	0.00	11.05	3.10	0.8	0.64	0.69	0.07	-0.66	0.73
S4	24.93	0.19	10.01	11.91	6.4	0.62	0.69	-0.08	-0.87	0.79
S5	47.93	0.00	22.65	15.58	5.6	0.20	0.72	-0.51	-0.98	0.46
S6	77.92	0.00	13.82	3.34	5.6	0.85	0.36	0.07	-0.49	0.56
S7	95.77	0.03	7.41	8.31	2.2	0.86	0.73	0.38	-0.70	1.08
S8	95.14	0.00	6.05	0.66	0.2	0.88	0.55	0.18	-0.76	0.95
S9	87.93	0.06	7.32	5.26	1.6	0.89	0.85	0.55	-0.83	1.38
S10	74.03	0.06	4.89	0.52	0.8	0.82	0.93	0.62	-0.61	1.23
S11	51.25	0.59	13.76	3.66	9	0.72	0.35	-0.27	-0.73	0.51
S12	62.76	0.00	6.88	0.72	2	0.85	0.75	0.42	-0.59	1.02

Descriptive statistics for the size of regions (after thresholding at $r = 0.35$) that correlate both positively and negatively with the PCC before global signal regression are shown for each subject. Corresponding sizes after global signal regression (i.e., the size of the task negative and task positive regions, respectively) are also given. The phase between average time series from the task-negative and task-positive regions before global signal regression is reported. Correlations between these two time series and the global signal and between each other are shown (GS/Neg, GS/Pos and Neg/Pos respectively). The correlation between the task-negative and task-positive time series (Neg/Pos) after global signal regression along with the difference in correlation value that the global signal caused are also reported. Definitions: +ve = positive, -ve = negative, PCC = posterior cingulate cortex, GS = global signal, Neg = task-negative time series, Pos = task-positive time series, GSR = global signal regression.

$CC = -0.5 \pm 0.13$ (Birn et al., 2006). One would expect the RVT technique to somewhat reduce correlation values in the task-positive network.

There also exists a large amount of variability across subjects in the correlation values in the task-positive areas before any correction is performed. Comparing these values with the histograms in Fig. 8 shows that subjects with the lowest average correlation values in task-positive areas (see Subjects 4 and 11) have distributions that extend the furthest into negative values. Voxels that become the most strongly anti-correlated after global signal regression are the voxels that were least correlated prior to the regression. This suggests that when global signal regression distorts the distribution to become bell-shaped and centred around zero (as shown experimentally), voxels in the task-positive network always remain on the negative tail. Depending on the shape of this distribution, their average correlation value before correction can be both positive and negative. After global signal regression, the variability in these values is removed since all distributions are forced to be bell-shaped (thus satisfying the theoretical prediction that correlation values sum to less than zero). This implies that the drastic reductions of the correlation values in the task-positive network are due to the mathematics of converting a skewed distribution into a bell-shaped distribution centred on zero. Task-positive areas may be converted into negative correlations simply because they are the least correlated with the PCC to begin with.

If the global signal obscures fluctuations related to neuronal firing, why does the regression of this signal not uncover other positively correlated areas? Large areas of the brain are significantly positively correlated with the PCC before removal of the global signal. The task-negative network produced after global signal regression is always a subset of these positive areas. The areas that subsequently comprise the task-negative network can largely be determined by raising the threshold on the correlation maps without global signal regression (see Fig. 6). The global correction does not reveal any new areas that are positively correlated with the PCC. If the neurons firing in the PCC produced fluctuations that were not correlated to the global confound but were overshadowed by such a confound, we would not expect the positively correlated areas to be identical before and after global signal regression. This suggests that either the global signal is highly correlated to the PCC neuronal fluctuations or that fluctuations due to neuronal firing in the PCC are visible above the global signal confound. If the former were true then neuronal fluctuations in the PCC would be removed by global signal regression. If the latter were true, then anti-correlated networks should be visible above the global signal confound also. Since the task-negative network remains highly correlated after global signal regression and anti-correlated networks are not visible

before, this suggests that a shifting of distributions by the technique is the most likely explanation for negatively correlated regions.

If negatively correlated voxels are an artefact of the correction technique, why are they in spatially contiguous areas and not randomly scattered throughout the brain? We have shown experimentally that global signal regression forces a skewed distribution to become bell-shaped, centred on zero and that voxels on the negative tail of this skewed distribution become the most negatively correlated. The reason that the task-positive network becomes anti-correlated after global signal regression is because it is the least correlated before. But why is it the least correlated? One possibility is that the correlation is reduced by a spatial heterogeneity in the “global” respiration related signal changes. Studies have shown that fMRI signal induced by variations in breathing depth (also reflected in variations in the end-tidal CO_2) have spatial structure and are most prominent in areas associated with the default mode network (Birn et al., 2006; Wise et al., 2004). When seed voxels are chosen from the PCC, an area in the default mode network that is highly contaminated with breathing rate related confounds, areas outside the default mode network such as the task-positive regions will likely be the least correlated with the seed voxels. Also, phase differences can exist between time series drawn from the task-negative and task-positive areas and that these time series can have a low correlation value between them even though both are highly correlated with the global signal (see Table 1). This could be due to differences in vascular supply or haemodynamic properties that are unrelated to neuronal firing. The phase simulations show that this technique can introduce negative correlations between time series that display identical resting state fluctuations that are slightly shifted in time. The largest phase change in these simulations corresponds to a time delay of between 1.25 s and 12.5 s for frequencies between 0.01 Hz and 0.1 Hz. Although values on the upper end of this range are unlikely to be accounted for by haemodynamic delays alone, values on the lower end are within range of natural haemodynamic delay variation (Aguirre et al., 1998b; Handwerker et al., 2004; Saad et al., 2001) suggesting that areas whose neurons fluctuate simultaneously but whose haemodynamics vary temporally could display anti-correlations after global signal regression. Furthermore, underlying neuronal delays in low-frequency fluctuations on the order of seconds between networks would display high correlations with each other neuronally but would be deemed anti-correlated after global signal regression. (It should also be noted that $\pi/4$ was an arbitrarily chosen phase delay and that if the maximum delay was smaller, similar correlation maps would be obtained). Spatial smoothing tends to increase the correlation of voxel

time series with the global signal (Aguirre et al., 1997) and could also help consolidate sparsely positioned negatively correlated voxels into contiguous regions.

If global signal regression is a poor choice of correction method, which techniques should be used to remove noise confounds, be they physiologically related or not? The circularity of choosing a noise regressor from the data itself leads to the problems identified in this study. Confounding regressors derived from external measures, such as pulse oximeter, respiration belt and end-tidal CO₂ readings, should not suffer from this problem. There exist other techniques that derive regressors from regions in the fMRI data, such as the ventricles, white matter, blood vessels and CSF, that may cause less of a problem. Studies have shown that anti-correlated networks do not exist after regression of movement, ventricle and white matter signals (Fox et al., 2008; Weissenbacher et al., 2008) and therefore these correction methods may not cause the same artefact as global signal regression. If a regressor is uncontaminated by the task in which we are interested (in this case the resting state fluctuations), then it may not introduce anti-correlations. However, the greater the correlation with the global signal, the larger a problem such a regression may cause by reducing true neuronally-related correlations and by increasing numbers of anti-correlated voxels. Using global signal regression to remove physiological noise clearly has confounding factors. While the optimal way to remove physiological noise isn't clear, many of the best current methods require collection of cardiac and respiration time courses. Collecting such data should become a standard part of resting state connectivity studies.

This study has concentrated on correlation-based analyses, the most popular method employed to detect resting state networks. Other methods have been used to support the idea that the brain is organized into independent neuronal networks, for example, *k*-means clustering (Golland et al., 2008) or ReHo analysis (Long et al., 2008). However, removing the global signal is often included as a pre-processing step. An alternative technique that is now frequently used to study functional connectivity is independent component analysis or ICA (Birn et al., 2008; Calhoun et al., 2005; Kiviniemi et al., 2003; Kiviniemi et al., 2004; Long et al., 2008; McKeown et al., 2003; McKeown et al., 1998; McKeown and Sejnowski, 1998). With ICA, statistically independent maps and their associated time courses are extracted from the data. If the global signal is not removed before this technique is employed, it must be assigned to one or more components. Therefore, ICA may have an inherent global signal regression built into it. The time course for each component determines the behaviour of the network beyond all other components, including those that represent the global signal. Therefore, task-negative and task-positive networks determined by ICA may have associated time courses that artificially appear to be anti-correlated.

Global signal regression not only introduces anti-correlated regions in the brain, it also reduces positive correlation values. If such a simple mathematical process used as a pre-processing step affects correlation values so drastically, can we trust correlation techniques to give a true measure of connectivity in the brain? What do correlation values actually mean? In a clinical context, do small differences in connectivity measures between patient groups and controls really reflect deficits in neuronal connectivity? More importantly, can the global signal regression introduce false differences between these two groups? These questions must be taken into account before we can truly understand resting state neuronal networks in the brain, their relationships to each other and their importance in clinical conditions.

Conclusions

In this paper we have shown using theory, simulations, task and resting state data that anti-correlated networks in the brain may be a consequence of the global signal regression pre-processing step.

Mathematically, this technique forces approximately half the voxels in the brain to become anti-correlated with a seed voxel. Global signal regression of data containing a known neuronal oscillation corrupted by a known respiration confound demonstrated that this technique is not successful in removing nuisance global effects and that the locations of anti-correlated areas are dependent on the relative phases of the global and seed voxel time series. In resting-state data, anti-correlated networks were not evident until global signal regression was performed. Other methods of low-frequency confound correction, such as RVT, have little effect on correlation values and, thus, do not reveal such networks. This study has demonstrated that employing global signal regression as a correction technique may cause spurious findings of negatively correlated regions in the brain. Great care should be exercised when interpreting such results.

Acknowledgments

This study was supported by the Intramural Research Program, National Institute of Mental Health, NIH. Thanks go to Phil Reiss and Clare Kelly for helpful discussions.

References

- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1997. Empirical analyses of BOLD fMRI statistics. II. Spatially smoothed data collected under null-hypothesis and experimental conditions. *Neuroimage* 5, 199–212.
- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998a. The inferential impact of global signal covariates in functional neuroimaging analyses. *Neuroimage* 8, 302–306.
- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998b. The variability of human, BOLD hemodynamic responses. *Neuroimage* 8, 360–369.
- Andersson, J.L., 1997. How to estimate global activity independent of changes in local activity. *Neuroimage* 6, 237–244.
- Andersson, J.L., Ashburner, J., Friston, K., 2001. A global estimator unbiased by local changes. *Neuroimage* 13, 1193–1206.
- Arndt, S., Cizadlo, T., O'Leary, D., Gold, S., Andreasen, N.C., 1996. Normalizing counts and cerebral blood flow intensity in functional imaging studies of the human brain. *Neuroimage* 3, 175–184.
- Beall, E.B., Lowe, M.J., 2007. Isolating physiologic noise sources with independently determined spatial measures. *Neuroimage* 37, 1286–1300.
- Birn, R.M., Diamond, J.B., Smith, M.A., Bandettini, P.A., 2006. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *Neuroimage* 31, 1536–1548.
- Birn, R.M., Murphy, K., Bandettini, P.A., 2008. The effect of respiration variations on independent component analysis results of resting state functional connectivity. *Hum. Brain Mapp.* 29, 740–750.
- Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34, 537–541.
- Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network: anatomy, function, and relevance to disease. *Ann. N. Y. Acad. Sci.* 1124, 1–38.
- Calhoun, V.D., Adali, T., Stevens, M.C., Kiehl, K.A., Pekar, J.J., 2005. Semi-blind ICA of fMRI: a method for utilizing hypothesis-derived time courses in a spatial ICA analysis. *Neuroimage* 25, 527–538.
- Chuang, K.H., Chen, J.H., 2001. IMPACT: image-based physiological artifacts estimation and correction technique for functional MRI. *Magn. Reson. Med.* 46, 344–353.
- Cohen, M.S., 1997. Parametric analysis of fMRI data using linear systems methods. *Neuroimage* 6, 93–103.
- Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173.
- Damoiseaux, J.S., Rombouts, S.A., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. *Proc. Natl. Acad. Sci. U. S. A.* 103, 13848–13853.
- De Luca, M., Beckmann, C.F., De Stefano, N., Matthews, P.M., Smith, S.M., 2006. fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage* 29, 1359–1367.
- Desjardins, A.E., Kiehl, K.A., Liddle, P.F., 2001. Removal of confounding effects of global signal in functional MRI analyses. *Neuroimage* 13, 751–758.
- Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711.
- Fox, P.T., Mintun, M.A., 1989. Noninvasive functional brain mapping by change-distribution analysis of averaged PET images of H215O tissue activity. *J. Nucl. Med.* 30, 141–149.
- Fox, P.T., Mintun, M.A., Reiman, E.M., Raichle, M.E., 1988. Enhanced detection of focal brain responses using intersubject averaging and change-distribution analysis of subtracted PET images. *J. Cereb. Blood Flow Metab.* 8, 642–653.
- Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 9673–9678.

- Fox, M.D., Snyder, A.Z., Zacks, J.M., Raichle, M.E., 2006. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nat. Neurosci.* 9, 23–25.
- Fox, M.D., Zhang, D., Snyder, A.Z., Raichle, M.E., 2008. Global signal regression and anticorrelations in resting state fMRI data. *Proc. HBM*, 575 W-AM.
- Fransson, P., 2005. Spontaneous low-frequency BOLD signal fluctuations: an fMRI investigation of the resting-state default mode of brain function hypothesis. *Hum. Brain Mapp.* 26, 15–29.
- Friston, K.J., Frith, C.D., Liddle, P.F., Dolan, R.J., Lammertsma, A.A., Frackowiak, R.S., 1990. The relationship between global and local changes in PET scans. *J. Cereb. Blood Flow Metab.* 10, 458–466.
- Garriety, A.G., Pearson, G.D., McKiernan, K., Lloyd, D., Kiehl, K.A., Calhoun, V.D., 2007. Aberrant “default mode” functional connectivity in schizophrenia. *Am. J. Psychiatry* 164, 450–457.
- Gavrilescu, M., Shaw, M.E., Stuart, G.W., Eckersley, P., Svalbe, I.D., Egan, G.F., 2002. Simulation of the effects of global normalization procedures in functional MRI. *Neuroimage* 17, 532–542.
- Glover, G.H., Li, T.Q., Ress, D., 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn. Reson. Med.* 44, 162–167.
- Golland, Y., Golland, P., Bentin, S., Malach, R., 2008. Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems. *Neuropsychologia* 46, 540–553.
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci. U. S. A.* 100, 253–258.
- Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc. Natl. Acad. Sci.* 101, 4637–4642.
- Handwerker, D.A., Ollinger, J.M., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* 21, 1639–1651.
- Hu, X., Le, T.H., Parrish, T., Erhard, P., 1995. Retrospective estimation and correction of physiological fluctuation in functional MRI. *Magn. Reson. Med.* 34, 201–212.
- Junghofer, M., Schupp, H.T., Stark, R., Vaitl, D., 2005. Neuroimaging of emotion: empirical effects of proportional global signal scaling in fMRI data analysis. *Neuroimage* 25, 520–526.
- Just, M.A., Cherkassky, V.L., Keller, T.A., Kana, R.K., Minshew, N.J., 2007. Functional and anatomical cortical underconnectivity in autism: evidence from an fMRI study of an executive function task and corpus callosum morphometry. *Cereb. Cortex* 17, 951–961.
- Kelly, A.M.C., Uddin, L.Q., Biswal, B.B., Castellanos, F.X., Milham, M.P., 2008. Competition between functional brain networks mediates behavioral variability. *Neuroimage* 39, 527–537.
- Kennedy, D.P., Redcay, E., Courchesne, E., 2006. Failing to deactivate: resting functional abnormalities in autism. *Proc. Natl. Acad. Sci. U. S. A.* 103, 8275–8280.
- Kiviniemi, V., Kantola, J.H., Jauhiainen, J., Hyvarinen, A., Tervonen, O., 2003. Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage* 19, 253–260.
- Kiviniemi, V., Kantola, J.H., Jauhiainen, J., Tervonen, O., 2004. Comparison of methods for detecting nondeterministic BOLD fluctuation in fMRI. *Magn. Reson. Imaging* 22, 197–203.
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., et al., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. U. S. A.* 89, 5675–5679.
- Li, T.Q., Kastrup, A., Takahashi, A.M., Moseley, M.E., 1999. Functional MRI of human brain during breath holding by BOLD and FAIR techniques. *Neuroimage* 9, 243–249.
- Long, X.Y., Zuo, X.N., Kiviniemi, V., Yang, Y., Zou, Q.H., Zhu, C.Z., Jiang, T.Z., Yang, H., Gong, Q.Y., Wang, L., Li, K.C., Xie, S., Zang, Y.F., 2008. Default mode network as revealed with multiple methods for resting-state functional MRI analysis. *J. Neurosci. Methods* 171, 349–355.
- Lowe, M.J., Mock, B.J., Sorenson, J.A., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. *Neuroimage* 7, 119–132.
- Lund, T.E., Madsen, K.H., Sidaros, K., Luo, W.L., Nichols, T.E., 2006. Non-white noise in fMRI: does modelling have an impact? *Neuroimage* 29, 54–66.
- Macey, P.M., Macey, K.E., Kumar, R., Harper, R.M., 2004. A method for removal of global effects from fMRI time series. *Neuroimage* 22, 360–366.
- Mason, M.F., Norton, M.I., Van Horn, J.D., Wegner, D.M., Grafton, S.T., Macrae, C.N., 2007. Wandering minds: the default network and stimulus-independent thought. *Science* 315, 393–395.
- McKeown, M.J., Sejnowski, T.J., 1998. Independent component analysis of fMRI data: examining the assumptions. *Hum. Brain Mapp.* 6, 368–372.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 6, 160–188.
- McKeown, M.J., Hansen, L.K., Sejnowski, T.J., 2003. Independent component analysis of functional MRI: what is signal and what is noise? *Curr. Opin. Neurobiol.* 13, 620–629.
- Perlberg, V., Bellec, P., Anton, J.L., Pelegrini-Issac, M., Doyon, J., Benali, H., 2007. CORSICA: correction of structured noise in fMRI by automatic identification of ICA components. *Magn. Reson. Imaging* 25, 35–46.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. *Proc. Natl. Acad. Sci.* 98, 676–682.
- Ramsay, S.C., Murphy, K., Shea, S.A., Friston, K.J., Lammertsma, A.A., Clark, J.C., Adams, L., Guz, A., Frackowiak, R.S., 1993. Changes in global cerebral blood flow in humans: effect on regional cerebral blood flow during a neural activation task. *J. Physiol.* 471, 521–534.
- Saad, Z.S., Ropella, K.M., Cox, R.W., DeYoe, E.A., 2001. Analysis and use of fMRI response delays. *Hum. Brain Mapp.* 13, 74–93.
- Shmueli, K., van Gelderen, P., de Zwart, J.A., Horowitz, S.G., Fukunaga, M., Jansma, J.M., Duyn, J.H., 2007. Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. *NeuroImage* 38, 306–320.
- Shulman, G.L., Corbetta, M., Buckner, R.L., Raichle, M.E., Fiez, J.A., Miezin, F.M., Petersen, S.E., 1997. Top-down modulation of early sensory cortex. *Cereb. Cortex* 7, 193–206.
- Sonuga-Barke, E.J., Castellanos, F.X., 2007. Spontaneous attentional fluctuations in impaired states and pathological conditions: a neurobiological hypothesis. *N. Biobehav. Rev.* 31, 977–986.
- Strother, S.C., Anderson, J.R., Schaper, K.A., Sidtis, J.J., Liow, J.S., Woods, R.P., Rottenberg, D.A., 1995. Principal component analysis and the scaled subprofile model compared to intersubject averaging and statistical parametric mapping: I. “Functional connectivity” of the human motor system studied with [¹⁵O]water PET. *J. Cereb. Blood Flow Metab.* 15, 738–753.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme Medical, New York.
- Vincent, J.L., Snyder, A.Z., Fox, M.D., Shannon, B.J., Andrews, J.R., Raichle, M.E., Buckner, R.L., 2006. Coherent spontaneous activity identifies a hippocampal–parietal memory network. *J. Neurophysiol.* 96, 3517–3531.
- Weissenbacher, A., Lanzenberger, R., Moser, E., Windischberger, C., 2008. Increasing specificity of resting-state fMRI-data using multiple regression analysis. *Proc. HBM*, TH-AM.
- Wise, R.G., Ide, K., Poulin, M.J., Tracey, I., 2004. Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal. *Neuroimage* 21, 1652–1664.
- Zarahn, E., Aguirre, G.K., D'Esposito, M., 1997. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage* 5, 179–197.